

## Data Mining Technologies for Blood Glucose and Diabetes Management

Riccardo Bellazzi, Ph.D.,<sup>1</sup> and Ameen Abu-Hanna, Ph.D.<sup>2</sup>

### Abstract

Data mining is the process of selecting, exploring, and modeling large amounts of data to discover unknown patterns or relationships useful to the data analyst. This article describes applications of data mining for the analysis of blood glucose and diabetes mellitus data. The diabetes management context is particularly well suited to a data mining approach. The availability of electronic health records and monitoring facilities, including telemedicine programs, is leading to accumulating huge data sets that are accessible to physicians, practitioners, and health care decision makers. Moreover, because diabetes is a lifelong disease, even data available for an individual patient may be massive and difficult to interpret. Finally, the capability of interpreting blood glucose readings is important not only in diabetes monitoring but also when monitoring patients in intensive care units. This article describes and illustrates work that has been carried out in our institutions in two areas in which data mining has a significant potential utility to researchers and clinical practitioners: analysis of (i) blood glucose home monitoring data of diabetes mellitus patients and (ii) blood glucose monitoring data from hospitalized intensive care unit patients.

*J Diabetes Sci Technol* 2009;3(3):603-612

### Introduction

Data mining is the process of selecting, exploring, and modeling large amounts of data to discover unknown patterns or relationships useful to the data analyst.<sup>1,2</sup> Notwithstanding the common understanding that data mining applies to data analysis problems with vast amounts of data and is solved by resorting to effective combination of analytical methods, the term data mining is still unspecific, largely due to its process-based nature. For this reason a large variety of methods

are considered part of data mining, including computer science approaches, such as multidimensional databases, machine learning, soft computing and data visualization, and statistical-based methods, including hypothesis testing, clustering, classification, and regression. Data mining technologies have proven to be useful and effective in different areas, including marketing, customer relationship management, engineering, and biomedical research.<sup>3,4</sup>

**Author Affiliations:** <sup>1</sup>Dipartimento di Informatica e Sistemistica, Università di Pavia, Pavia, Italy; and <sup>2</sup>Department of Medical Informatics, University of Amsterdam, Amsterdam, The Netherlands

**Abbreviations:** (BG) blood glucose, (BGL) blood glucose level, (DM) diabetes mellitus, (ICU) intensive care unit, (IIT) insulin-intensive therapy, (PRIM) patient rule induction method, (TA) temporal abstraction

**Keywords:** blood glucose management, diabetes mellitus, data analysis, data mining

**Corresponding Author:** Riccardo Bellazzi, Ph.D., Dipartimento di Informatica e Sistemistica, Università di Pavia, Via Ferrata 1, 27100 Pavia, Italy; email address [riccardo.bellazzi@unipv.it](mailto:riccardo.bellazzi@unipv.it)

The goals of data mining can be classified into two tasks: description and prediction.<sup>3</sup> While the purpose of description is to extract understandable patterns and associations from data, the goal of prediction is to forecast one or more variables of interest. The main difference between the two tasks is therefore the presence of a response variable in the case of prediction problems: if the response variable is continuous, such as blood glucose (BG), the prediction task is said to be a regression problem, whereas if the response variable is categorical, such as hypertension (taking the values yes or no), the task is said to be a classification problem.

This article describes applications of data mining to medicine, particularly the analysis of diabetes mellitus (DM) data. Over the last years several papers have been published on the problem of analyzing DM data in the computer science community. In 1994, for example, the AAAI spring symposium on artificial intelligence in medicine published a BG home monitoring data set as a challenge for applying machine learning and artificial intelligence methods to describe and predict BG values.<sup>5</sup> Another well-known data set available at the University of California Irvine data repository<sup>6</sup> is the so-called “Pima Indians” database, which is a collection of medical diagnostic reports of 768 examples from a native American population living near Phoenix, Arizona. This data set was used to develop algorithms to predict whether a patient shows signs of diabetes according to 1990 World Health Organization criteria. Less frequent has been the application of data mining to real clinical problems, although some work has been carried out in the late 1990s, which were motivated primarily by telemedicine projects.

However, the DM context seems particularly well suited to a data mining approach. First of all, thanks to the current availability of electronic health records, administrative health records, and monitoring facilities, including telemedicine programs, the sheer size of data accumulated and accessible to physicians, practitioners, and health care decision makers is huge. Second, because DM is a lifelong disease, even data available for an individual patient turn out to be massive and difficult to interpret. Moreover, because the treatment of DM is multifaceted, the need for taking into account many variables is becoming compelling. Finally, the capability of interpreting BG readings is important not only in DM but also when monitoring patients in intensive care units (ICU). In particular, hyperglycemia is frequently encountered in critically ill patients, and evidence shows that normoglycemia (often defined as BG in the range

of 4.4–6.1 mmol/liter, or equivalently 80–110 mg/dl) can decrease the morbidity and mortality of such patients.<sup>7</sup> However, the implementation of glucose control with insulin clearly comes with the risk of severe hypoglycemia. The how and for whom glucose control is safe and effective remain quite elusive.<sup>8</sup> Data mining can help address these questions by focusing on patient subgroups whose BG over time exhibits markedly different behavior than the rest. In addition, clinicians and researchers are unsure about which parameters (mean BG, time to capture the normoglycemia range, etc.) indicate good quality of the BG regulatory process itself.<sup>9</sup> Data mining can be used to investigate which indicators approximate the clinicians’ perception of a good stable regulatory process.

For these reasons, this article describes the work that has been carried out by the authors and their collaborators in two areas in which data mining can be helpful to researchers and clinical practitioners: the analysis of (i) BG home monitoring data of DM patients and (ii) BG monitoring data from hospitalized ICU patients.

## Data Mining Techniques for Interpreting Blood Glucose Self-Monitoring Data

Since the early 1980s a variety of methods and tools have been designed and implemented to support health care givers and patients in interpreting BG home monitoring data. Some of the proposed approaches are now included in commercial tools applied routinely during home monitoring periodic data review. The usual way to review the outcome of a certain therapy scheme starting from the analysis of BG data is to check if they show a *cyclostationary pattern*, i.e., if the BG daily behavior is approximately the same over the monitoring time. A cyclostationary behavior is therefore characterized by the absence of significant trends (stationarity) and by a periodic (with period equal to 1 day) behavior of BG. The characteristic daily BG pattern that summarizes the typical patient’s response to the therapy is called “modal day” and is visualized in different ways: it may be shown by plotting all BG data on a 24-hour scale or by computing the frequency histograms of BG measurements at different times of the day, which are also called time slices.<sup>10</sup> Examples of time slices are “before breakfast,” “after breakfast,” “before lunch,” “after lunch,” and so on. This way of showing data is certainly useful, as it synthesizes the BG time series “at a glance.” If combined with lifestyle information, such as mealtimes, it can be very helpful in highlighting the main components of the BG profiles.

However, the modal day visualization has several limitations—the main one being that the BG time series is often neither cyclic nor stationary. The stochastic nature of the BG measurements and the day-by-day lifestyle and meal variation of the patient's behavior may lead to confusing results when all data are grouped in a single picture. For this reason, we have proposed<sup>11–13</sup> the combination of signal processing and artificial intelligence techniques to provide an alternative representation of BG data, aimed at discerning the modal day representation.

First of all, following the work of Deutsch and colleagues,<sup>14</sup> we applied a filtering approach, known as structural analysis, to search for prototypical structures in data. Structural analysis decomposes a time series into a collection of signals with known “temporal structure.” In more detail, the basic assumption of the method is that each BG measurement can be expressed as a sum of separate components: a trend component ( $T$ ), a cyclic component ( $C$ ), and a stochastic component ( $\epsilon$ ) so that, for each measurement,  $i$ :

$$BG_i = T_i + C_i + \epsilon_i. \quad (1)$$

The problem of extracting  $T_i$  and  $C_i$  can be solved mathematically in different ways, including Kalman filtering, least-squares fitting, and Bayesian smoothing.<sup>12</sup> By using the same technique it is possible to also separate daily cycles from weekly or monthly cycles, which may correspond to changes in the lifestyle during the weekend or through the year.

In our case, we represented the trend component using a random walk model, and the cyclic component as the composition of sine and cosine waves, in order to seek daily cycles.

Given the availability of trend and cyclic components, it is possible to pose some interesting clinical questions: (1) is there a trend in BG data, i.e., increasing, decreasing, or even oscillating periods, which may last for several monitoring days? (2) Is there one or more cyclic behavior in BG data? Even more interestingly, it is possible to combine trend and cycle information with the absolute value of BG, asking, for example, whether there are episodes in which there was a persistent hyperglycemia with a decreasing trend.

Following a data mining approach, a proper way to answer the aforementioned questions directly is to apply an artificial intelligence technique known as temporal abstractions (TAs).<sup>11</sup> The principle of TAs is to provide

an interval-based representation of monitoring data: time-stamped data are aggregated into time intervals during which a certain event occurs. In the case of BG monitoring data it seems natural to apply three different kinds of TAs: state TAs (e.g., “high” BG values), trends TAs (e.g., “increasing” BG values), and complex TAs (e.g., a daily pattern: “high in the morning and low at dinner”).

State TA detection corresponds to the search for periods in which BG was persistently in a clinical significant range: in our case we limited “abstractions” on the BG value to five possible qualitative levels—very high, high, normal, low, or very low. Such abstractions may correspond to fixed or patient-tailored threshold values for BG.<sup>11</sup>

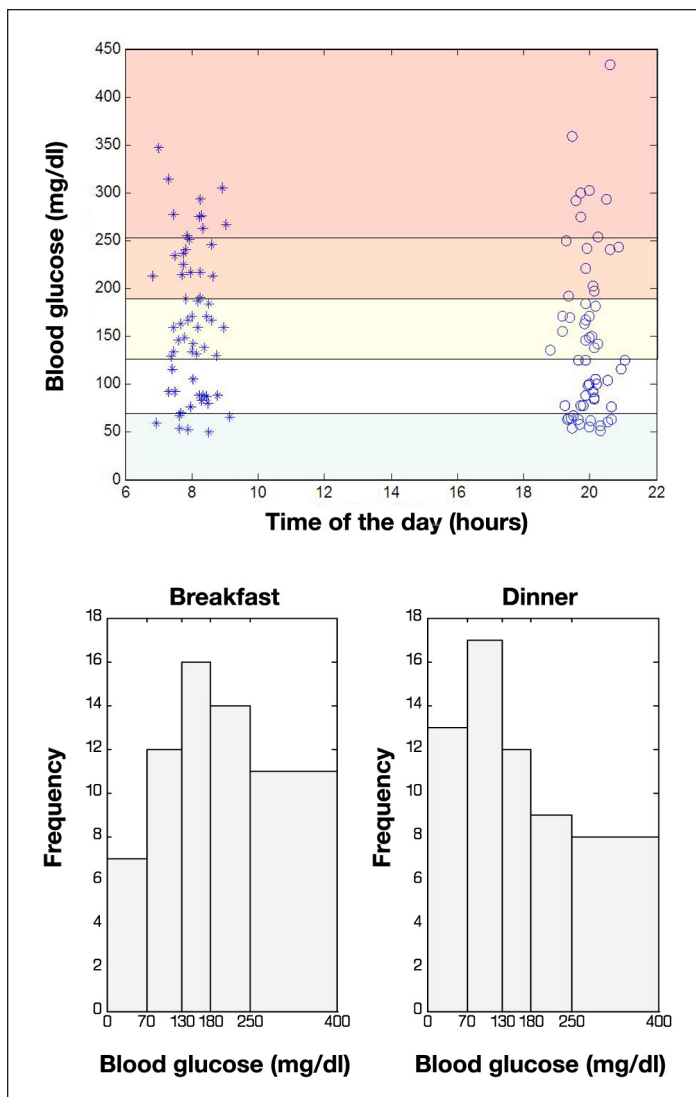
In terms of trend TA detection, we applied the trend TA method (for details, see Bellazzi and colleagues<sup>11</sup>) to the  $T$  time series obtained by decomposing the BG time series with **Equation (1)**. In this way it is possible to highlight the periods in which BG was increasing or decreasing, and it is easy to calculate some simple statistics on the different trend prevalence values over a certain monitoring period.

Finally, we applied a complex TA search on the  $C$  component of **Equation (1)**: every day was analyzed and the list of daily time measurements was sorted (e.g., if, given three measurements per day, the maximum BG measurement is at lunch and the minimum is at breakfast, then the pattern is < lunch, dinner, breakfast >). Days with the same BG cyclic patterns are then aggregated in order to check the persistence of that pattern over different days.

The combination of the three different pattern search mechanisms may be therefore used effectively to show a more complete picture of the behavior of a certain patient over time. **Figures 1** and **2** show an example of the application of the proposed method to a BG time series.

**Figure 1** reports on a modal day computed on 60 monitoring days for an 11-year-old male type 1 DM patient. Looking at these data it seems that BG data are distributed evenly over all five clinical intervals, with hypoglycemia mainly concentrated at dinner time.

**Figure 2** shows that analysis of the BG time series with structural decomposition and temporal abstractions highlights different behaviors in the control period: several periods of increasing and decreasing trends are



**Figure 1.** Modal day analysis (**top**) applied to BG time series coming from 60 days of home monitoring for an 11-year-old type 1 diabetic patient. Data are divided into five levels, including hypoglycemia, normal values, and three levels of hyperglycemia, from moderate to severe. Many measurements collected around 8 p.m. are hypoglycemia. Histograms of BG data (**bottom**) confirm that dinner time is critical for hypoglycemia, whereas higher values are collected at breakfast.

observed and a stable cyclical pattern of BG behavior is discerned from the 20th to the 30th monitoring day, corresponding to the 40 and 60 monitoring points, during which the measurement at dinner time was always lower than the one at breakfast time.

Interestingly, the application of TA methods to the analysis of BG data can be seen as a way to search for patterns in continuous monitoring data. For example, it is easy to find the pattern of metabolic instability, defined as a sequence of increasing and decreasing trends.

Such a pattern is searched by (i) detecting all occurrences of “BG increase” and “BG decrease” and (ii) extracting the intervals in which a “BG increase” is immediately followed by a “BG decrease” or vice versa.

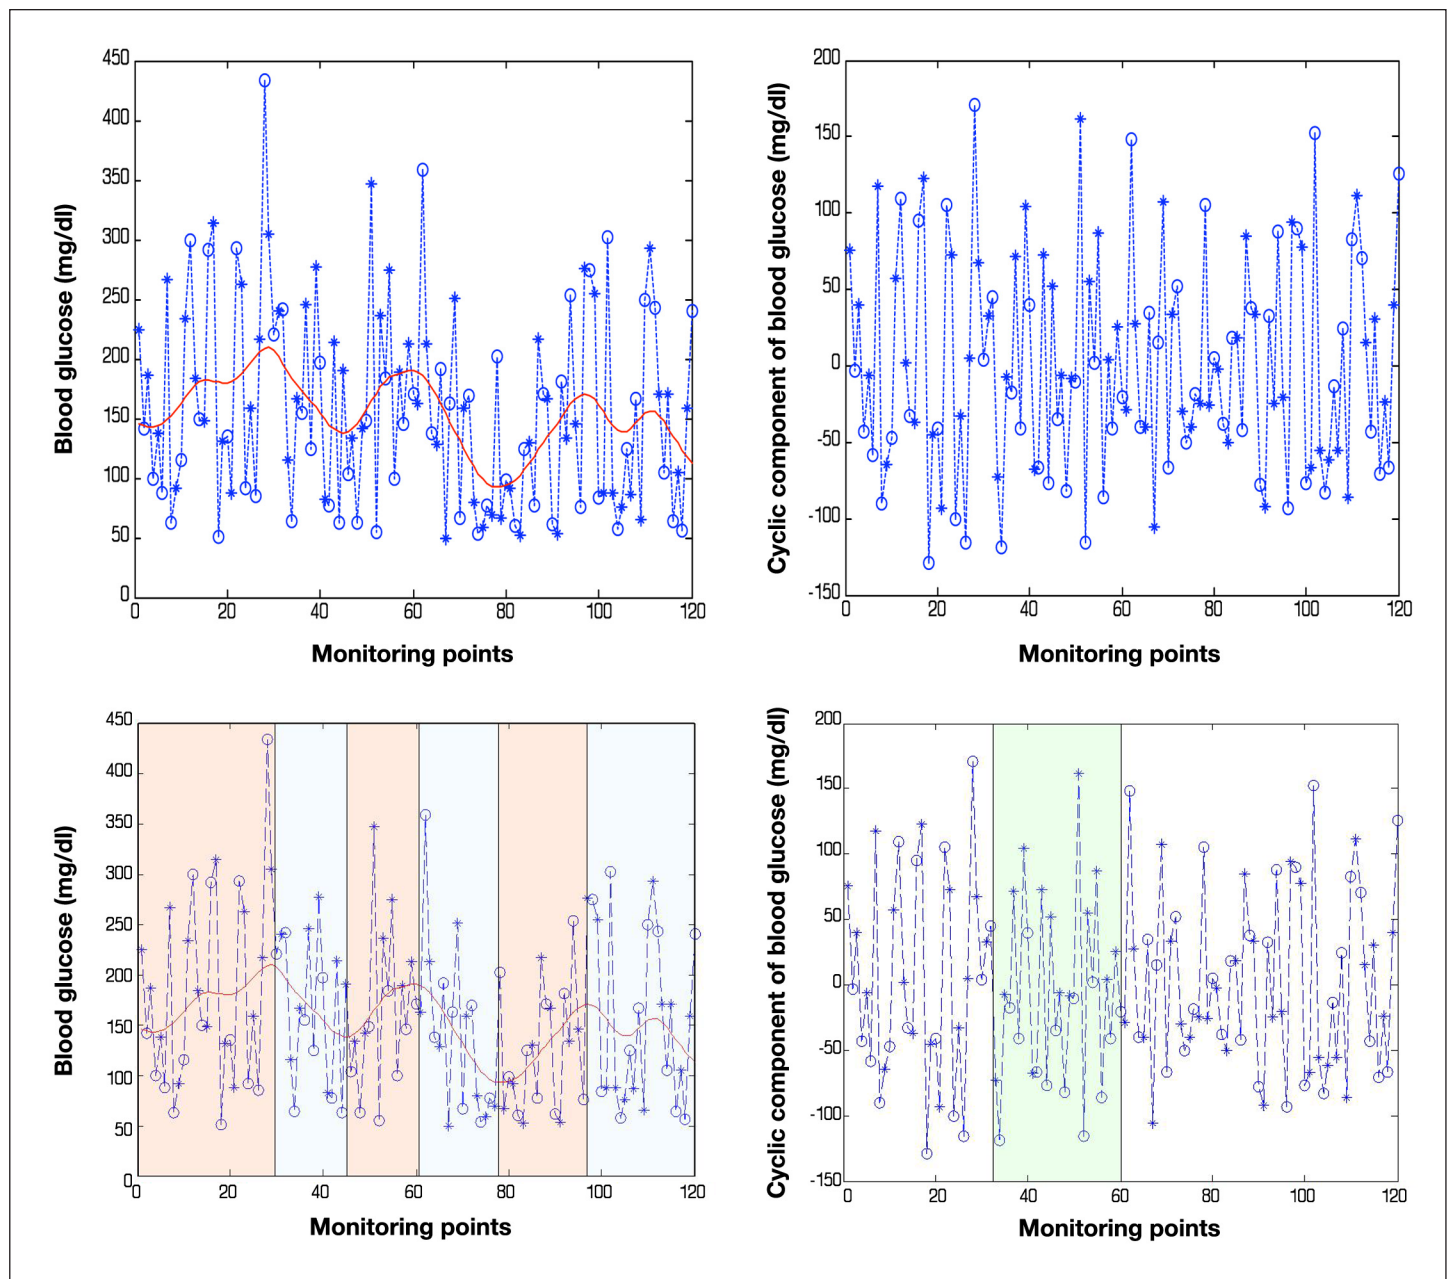
Bellazzi and colleagues<sup>11</sup> applied TA methods to derive a summary of the patients’ metabolic control over the monitoring periods between two face-to-face visits. Every period was characterized by TA-based summaries for the hyperglycemia, hypoglycemia, and metabolic instability episodes; summaries were derived by simply counting the total time span of the different episodes and by dividing such an extent by the monitoring time span.

It is important to note that the methods presented are able to perform an automatic “intelligent” analysis of blood glucose level (BGL) data, searching for interesting patterns by mimicking the data analysis strategies suggested by clinical experience. Such methods serve two main purposes: (a) to implement and standardize reasoning strategies about the single patient’s data and (b) to retrospectively extract information from a set of patients, thus highlighting the most frequent patterns and/or finding new patients subgroups.

## Data Mining Techniques for Interpreting ICU BG data

Glucose regulation is an increasingly important topic in intensive care where ways for improving guidelines to manage the BG of patients are constantly sought. For a long time it was common practice to administer insulin only when blood glucose levels exceeded 10–11.1 mmol/liter (180–200 mg/dl), as hyperglycemia was considered an adaptive response to critical illness. However, the landmark study by van den Berghe and colleagues<sup>7</sup> showed that normalization of the plasma glucose level of intensive care patients resulted in decreased morbidity and mortality. Notwithstanding the debate about the reproducibility of similar results in more recent trials,<sup>15</sup> new guidelines for strict BG regulation have emerged and are being widely adopted. These insulin-intensive therapy (IIT) guidelines strive to keep the BG in a strict range, such as between 4.4 and 6.1 mmol/liter (80–110 mg/dl). Guidelines are, however, not always beneficial to all patients at all times, and providing tools to investigate the effects of guideline-based therapy on clinical outcomes is an important contribution toward improving the guidelines. In general, there are three classes of clinically relevant questions one may want to address.





**Figure 2.** Structural analysis and temporal abstractions applied to the blood glucose time series reported in **Figure 1**. The patient measured BG values only at breakfast (stars) and dinner (circle) times. Once the trend time series is extracted from data (**upper left**), by applying trend TAs, it is easy to detect increasing (red box) and decreasing (cyan box) periods (**lower left**). Moreover, starting from the cycle time series (**upper right**), it is possible to highlight periods in which the daily pattern was the same. In the example the maximum was at breakfast and the minimum was at dinner (green box, **lower right**). Such information complements the statistics that can be obtained by interpreting **Figure 1**.

1. *Patients at risk.* Although these guidelines have the potential to lower the mean blood glucose level of the patient population as a whole, hyperglycemia is still often found in various critically ill patients. In addition, implementation of these IIT guidelines is nearly always associated with a significant increase in hypoglycemia.<sup>9</sup> One is interested in knowing “which patients are at high risk of hyperglycemia despite having an IIT guideline in

place” and “which patients are at very high risk of severe hypoglycemia because of these guidelines?”

2. *Role of last glucose measurement.* The great majority of current IIT guidelines use solely the last measured BG value of the patient in order to decide on the amount of insulin to be administered. Is there justification for this choice or should other factors be considered?

3. *Quality of the regulatory process.* In a recent systematic review<sup>9</sup> we encountered 30 different indicators used in studies for measuring the quality of the BG regulatory process. These indicators fell under the following categories: blood glucose zones (e.g., “hypoglycemia”), blood glucose levels (e.g., “mean blood glucose level”), time intervals (e.g., “time to occurrence of an event”), and protocol characteristics (e.g., “blood glucose sampling frequency”). Although each indicator provides some insight about the quality of the regulatory process, we clearly do not have a good understanding of what really comprises a good regulatory process. This wide variety of indicators hinders the comparability of studies.

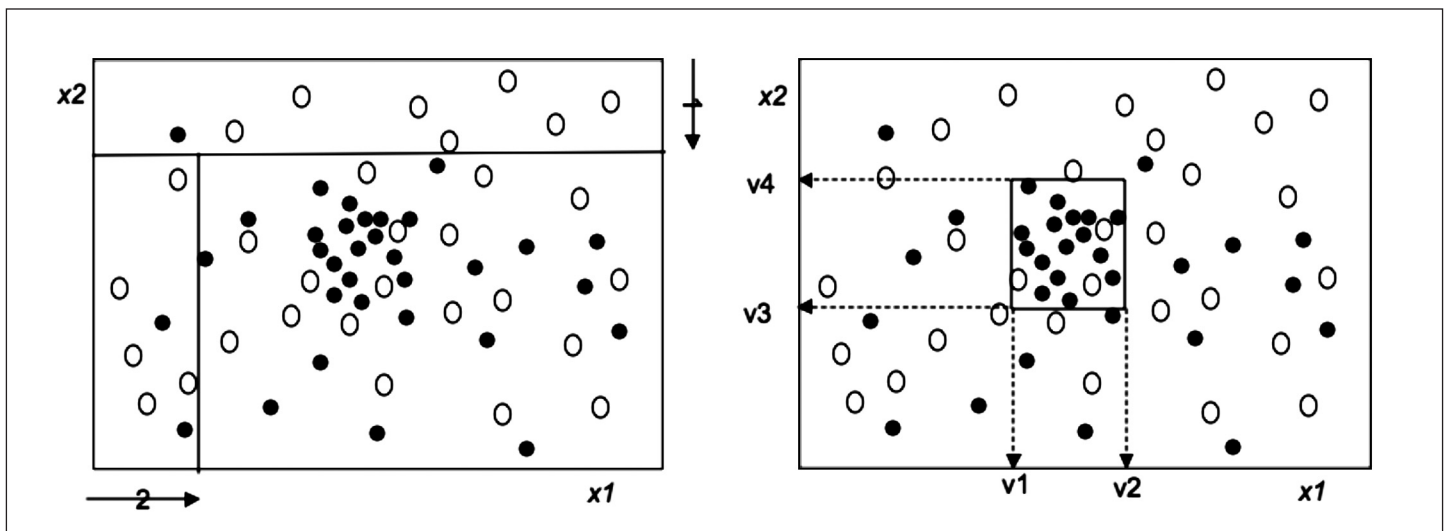
Data mining has a role to play in each of these problem classes.

### Patients at Risk

One approach to find patients at risk is to develop a model to directly predict the glucose level for any patient at any time based on modeling the underlying biological process and the insulin resistance dynamics themselves, as was attempted elsewhere.<sup>16</sup> Although this approach may provide insight in the underlying biological process, the complexity of this process (influenced by the many factors involved and their intricate interaction) impedes the development of models with sufficiently reliable predictions. An alternative approach that we pursued is focusing attention on observations that deviate markedly

from the “rest” of the observations. This concept has been referred to in the data mining and machine learning literature as subgroup discovery<sup>17,18</sup> and in statistical machine learning as bump hunting.<sup>19</sup> We hence do not seek a model that predicts the corresponding BG for any circumstance but instead only seek regions in feature space (the space that defines the input variables) for which we can predict significantly high (or low) blood glucose levels.

Let us illustrate this concept for finding subgroups (or “bumps” in feature space) with markedly high BG from recent work.<sup>20</sup> To this end, we have employed the patient rule induction method (PRIM) proposed by Friedman and Fisher<sup>19</sup> to routinely collected data. It is easy to illustrate how PRIM operates with two input variables, although the algorithm is meant to work with high-dimensional data. To simplify the illustration further, let us consider only two categories of BG values: high and low. In addition, let us ignore for the moment categorical input variables. Now consider **Figure 3**, which depicts high and low BG observations (dark and hollow circles, respectively) in a two-dimensional space defined by the variables  $x1$  and  $x2$ . At the outset, PRIM includes all sample observations in the initial “box.” It then attempts to shrink the box iteratively by peeling off a user-specified proportion ( $\alpha$ ) of data at the  $\alpha$  and  $1 - \alpha$  quantiles of a variable, such that there is maximal increase in the mean at each successive subbox.



**Figure 3.** Steps taken by PRIM to discover a subgroup with high values of BG. Solid circles indicate high values of BG and hollow ones indicate low BG values. The figure on the left-hand side shows the first two steps in the algorithm. In the first step, observations with the highest proportion  $\alpha$  of variable  $x2$  are removed. In the second step, the lowest proportion  $\alpha$  of the remaining observations from variable  $x1$  are removed. The figure on the right-hand side shows the final subgroup, defined here by observations with  $v1 < x1 < v2$  and  $v3 < x2 < v4$ . The proportion of observations in the subgroup should be at least  $\beta$  of the total sample. The meta-parameters  $\alpha$  (peeling fraction) and  $\beta$  (support) control the induction process.

“Peeling” follows essentially a hill-climbing search strategy in which each variable is considered in isolation. This peeling process continues by removing the proportion  $\alpha$  of the remaining observations until a user-specified minimum proportion ( $\beta$ ) of the initial sample in the box is reached. For a categorical variable, PRIM inspects the removal of observations belonging to each one of the possible categories separately.

When a subgroup has been found, one may continue seeking additional subgroups by removing the observations in the last discovered subgroup and repeating the process with the rest of the observations. The term “patient” in the algorithm refers to the fact that “peeling” removes only a small proportion of the observations in each step, unlike more greedy approaches, such as decision or regression trees.

A discovered subgroup is described by an *if-then* rule. The condition consists of the conjunctive term describing the hyperrectangle bounding the subgroup such as the rectangle defined by “ $v1 < x1 < v2$  and  $v3 < x2 < v4$ ” in **Figure 3**. The conclusion consists of a summary measure (usually the mean) of the response variable, here BGL, of the subgroup. An example of a rule discovered in our study<sup>20</sup> is shown in **Table 1**. For this subgroup, the mean BGL was 9.1 mmol/liter (=163.8 mg/dl).

As this rule reveals, the variables used include temporal abstractions, such as the most recent value (or the mean) during a time interval, of a series of values observed prior to the prediction. The decision on which abstractions to use was determined by a clinical expert. The algorithm discovers which abstractions and patterns are actually useful. From the rule described in **Table 1**

one uncovers, for example, that low body temperature and low bicarbonate for medically admitted patients are associated with hyperglycemia. This concords with expert knowledge, for example, intensive care specialists notice that patients with low body temperature do not react well to insulin therapy. It is important to note that such patterns are difficult to obtain from a traditional representation of the blood glucose values in the patient record. To address the question of which patients are at risk of hypoglycemia, instead of hyperglycemia, one seeks regions with very low values of BGL. Technically, one only needs to take the negative sign of the BG values and still run the same algorithm to seek regions with high BGLs (in this case those closest to 0).

### Role of Last Glucose Measurement

We now address the question on justification of the dominance of the last glucose value used in IIT guidelines to determine the amount of the insulin to be administered. One only needs to add BG values to routinely collected data and then simply perform the same subgroup analysis described earlier. In our application of this approach we obtained the following first rule<sup>20</sup>:

**if**      *the previous glucose measurement >13.2 mmol/liter*  
(238 mg/dl)

**and**    *the most recent bicarbonate measurement during the last 6 hours <26 mmol/liter*

**then**    *mean BG = 15.1 mmol/liter (272 mg/dl).*

In the discovery of subsequent subgroups the previous glucose measurement was always a dominant part of the condition of the corresponding rule. This finding provides a justification for use of the previous BG value to steer therapy in IIT guidelines, although it opens up the possibility to include other variables.

### Quality of the Regulatory Process

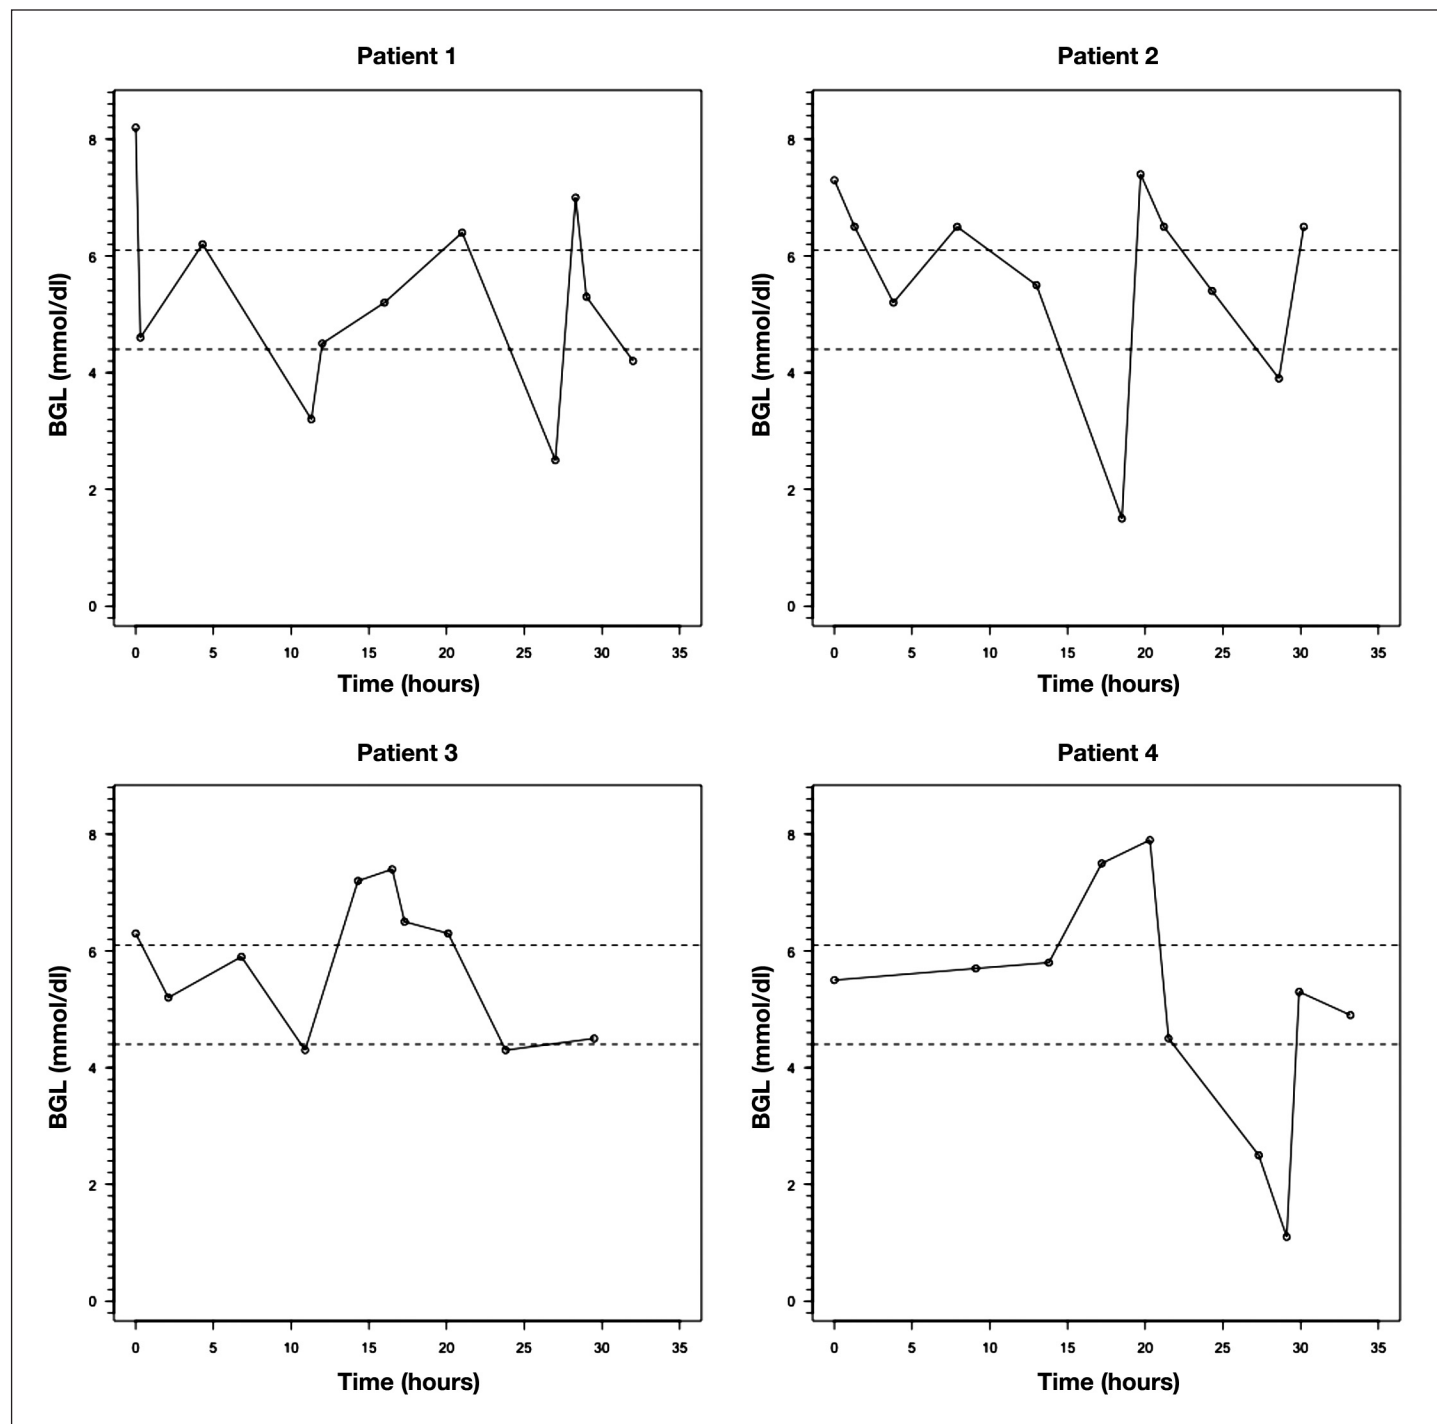
There is no agreement about what constitutes an appropriate indicator for the quality of the regulatory process. Consider **Figure 4**, which shows the BG for four patients over time. Clearly, simple indicators such as mean BG mask important details of the process, as they allow very high BG values to compensate very low values. In fact, the mean BG of each of the four patients in **Figure 4** falls in the BG target range. We suggested<sup>9</sup> an empirical approach to this problem. In particular, we can elicit the perception of a committee

**Table 1.**  
**Definition of a Discovered Subgroup with High Blood Glucose Values**

Condition	Variable description
Bicarbonate <20.5 mmol/liter	The <b>most recent</b> bicarbonate measurement <b>during the last 6 hours</b>
Admission type = medical	The medical reason for admission
Urine <2 liter or urine >4 liter	The amount of urine after 24 hours
21.5 < albumin <38.5 g/liter	The <b>most recent</b> albumin <b>during the last 24 hours</b>
Temperature <36.85°C	The <b>mean</b> body temperature <b>during the last 6 hours</b>

of experts (intensivists) of the quality of the processes for specific patients. For example, the committee can be shown BG courses over time of a set of a hundred patients. For each BG course a (qualitative, rank, or quantitative) quality score is given. Now the problem has been transformed to a supervised problem and hence we

can use data mining to approximate the scoring function of the experts. The input variables will consist of current and perhaps newly designed quality indicators, and the output variable is the experts' score. We hence seek to find the combination of indicators that best predicts this score.



**Figure 4.** The course of BG over time for four patients. Dashed horizontal lines indicate the required BG target range. Simple indicators, such as the mean BG per patient (respectively, 5.2, 5.7, 5.8, and 5.1 mmol/dl), do not seem to capture the quality of the regulatory process and a combination of such indicators should be sought.



## Discussion and Conclusions

Data mining technologies can have important utility in diabetes mellitus and blood glucose management. The increasing availability of data in an electronic format and the increase in the “intensity” of data collection made available by continuous glucose monitors and by telemedicine and pervasive computing<sup>21</sup> are driving forces for data mining in diabetes. As shown in this review of our own work, a data mining approach can be applied successfully in BG data analysis, both when data come from home monitoring of diabetic patients and when they are collected in intensive care units. In the future, we plan to demonstrate the usefulness of this kind of study by measuring the extent to which data mining approaches empower clinical research and practice.

Interestingly, such methods can also be applied to deal with the analysis of various types of data collected by health care organizations when performing their day-to-day activities. These data may be referred to as *process data*, as they carry information about the general behavior of the health care process. Process data may include administrative information, such as cost claims, physicians’ workload, and patients’ hospital admissions, as well as clinical variables, such as drug prescriptions and laboratory test results. A promising research and application area is therefore the exploitation of process data to highlight critical behavior of the health care providers in order to guide them following an “organizational learning strategy.”<sup>22,23</sup> In this field we are collaborating with health care agencies to derive “rules” that describe the most frequent temporal patterns of diabetic outpatients in terms of sequences of drug prescriptions, prescribed laboratory test, visits, and hospital admissions.<sup>24</sup>

### Acknowledgments:

We gratefully acknowledge the contribution of all researchers who collaborated in the work reviewed in this article: Giuseppe d’Annunzio, Giuseppe De Nicolao, Cristiana Larizza, Stefania Montani, Paolo Magni, Mario Stefanelli, Rob Bosman, Saeid Eslami, Evert de Jonge, Nicolette de Keizer, Barry Nannings, and Marcus Schultz. The work presented was supported by the M<sup>2</sup>DM project, funded by the European Commission, and by the Netherlands Organization for Scientific Research (NWO) under the I-Catcher project, Number 634.000.020.

### References:

- Giudici P. Applied data mining, statistical methods for business and industry. Wiley & Sons; 2003.
- Fayyad U, Piatetsky-Shapiro G, Smyth P. Data mining and knowledge discovery in databases. *Commun ACM*. 1996;39:24-6.
- Bellazzi R, Zupan B. Predictive data mining in clinical medicine: current issues and guidelines. *Int J Med Inform*. 2008;77(2):81-97.
- Bellazzi R, Zupan B. Towards knowledge-based gene expression data mining. *J Biomed Inform*. 2007;40(6):787-802.
- Kohane I, Uckun S. Artificial intelligence in medicine: interpreting clinical data. AAAI technical report SS-94-01 [cited 2008 Dec 14]. Available from: <http://www.aaai.org/Press/Reports/Symposia/Spring/spring-reports.php#SS94>.
- UCI machine learning repository [cited 2008 Dec 14]. Available from: <http://archive.ics.uci.edu/ml/>.
- Van den Berghe G, Wouters P, Weekers F, Verwaest C, Bruyninckx F, Schetz M, Vlasselaers D, Ferdinande P, Lauwers P, Bouillon R. Intensive insulin therapy in critically ill patients. *N Engl J Med*. 2001;345(19):1359-67.
- Chase JG, Shaw GM. Is there more to glycaemic control than meets the eye? *Crit Care*. 2007;11(4):160.
- Eslami S, de Keizer NF, de Jonge E, Schultz MJ, Abu-Hanna A. A systematic review on quality indicators for tight glycaemic control in critically ill patients: need for an unambiguous indicator reference subset. *Crit Care*. 2008;11(12):R139.
- Lehmann ED. Application of information technology in clinical diabetes care--a special issue. Part 1. Databases, algorithms and decision support. *Med Inform (Lond)*. 1996;21(4):255-8.
- Bellazzi R, Larizza C, Magni P, Montani S, Stefanelli M. Intelligent analysis of clinical time series: an application in the diabetes mellitus domain. *Artif Intell Med*. 2000;20(1):37-57.
- Bellazzi R, Magni P, De Nicolao G. Bayesian analysis of blood glucose time series from diabetes home monitoring. *IEEE Trans Biomed Eng*. 2000;47(7):971-5.
- Magni P, Bellazzi R. A stochastic model to assess the variability of blood glucose time series in diabetic patients self-monitoring. *IEEE Trans Biomed Eng*. 2006;53(6):977-85.
- Deutsch T, Lehmann ED, Carson ER, Roudsari AV, Hopkins KD, Sönksen PH. Time series analysis and control of blood glucose levels in diabetic patients. *Comput Methods Programs Biomed*. 1994;41(3-4):167-82.
- Schultz MJ, Spronk PE, Abu-Hanna A, Afessa B, Gajic O. Evidence-lost to tight glycemic control? *Crit Care*. 2008;12(5):430.
- Van Herpe T, Espinoza M, Haverbeke N, De Moor B. Glycemia prediction in critically ill patients using an adaptive modeling approach. *J Diabetes Sci Technol*. 2007;1(3):348-56.
- Siebes AP. Data surveyor. In Kloesgen W, Zytkow JM, editors. *Handbook of data mining and knowledge discovery*. Oxford: Oxford University Press; 2002. p. 572-75.
- Lavrac N, Kavsek B, Flach PA, Todorovski L. Subgroup discovery with CN2-SD. *J Mach Learn Res*. 2004;5:153-88.
- Friedman JH, Fisher NI. Bump hunting in high-dimensional data (with discussion). *Stat Comput*. 1999;9:123-62.
- Nannings B, Bosman RJ, Abu-Hanna A. A subgroup discovery approach for scrutinizing blood glucose management guidelines by the identification of hyperglycemia determinants in ICU patients. *Methods Inf Med*. 2008;47(6):480-8.
- Bellazzi R. Telemedicine and diabetes management: current challenges and future research directions. *J Diabetes Sci Technol*. 2008;2(1):98-104.

22. Stefanelli M. The socio-organizational age of artificial intelligence in medicine. *Artif Intell Med.* 2001;23(1):25-47.
23. Mullins IM, Siadaty MS, Lyman J, Scully K, Garrett CT, Miller WG, Muller R, Robson B, Apte C, Weiss S, Rigoutsos I, Platt D, Cohen S, Knaus WA. Data mining and clinical data repositories: Insights from a 667,000 patient data set. *Comput Biol Med.* 2006;36(12):1351-77.
24. Concaro S, Sacchi L, Cerra C, Fratino P, Bellazzi R. Temporal data mining for the analysis of administrative healthcare data. IDAMAP 2008 working notes. Available from: <http://www.idamap.org/>.